

A Journey Instead of a Calculation

Imagine someone asks you to drive from one side of a large city to the other.

You don't sit in the driveway calculating every turn of the steering wheel before you leave. That would be impossible.

Instead, you start driving.

You look ahead.

You make a tiny steering correction.

You move forward a few metres.

You look again.

Another small correction.

Thousands of these tiny adjustments eventually get you to your destination.

Now imagine that a neural network works in much the same way.

People often picture a neural network as performing one enormous, mysterious calculation. But another way to think about it is as a long journey made up of many small decisions.

Each layer asks a simple question:

"Given where I am now, what's the next small adjustment?"

It doesn't need to solve the whole problem. It only needs to make things a little better for the next layer.

If one layer doesn't get things exactly right, that's fine. The next few layers can compensate, just as a driver gently corrects the car after drifting slightly off the ideal line.

The really interesting part is how each layer decides what kind of adjustment to make.

Imagine you're driving. Your steering depends entirely on the situation around you.

A straight motorway calls for one kind of steering.

A winding country lane calls for another.

A busy roundabout calls for something different again.

You don't use the same driving rule everywhere. The local context determines what you do next.

Many modern neural networks behave in a similar way.

Rather than always using exactly the same computation, they first look at the current situation. That context determines which learned pieces of computation should become active. It's a bit like opening a toolbox and choosing the tools that fit the job at that particular moment.

The network has learned many useful computational fragments during training. The context simply decides which ones are most appropriate right now.

This is why deep networks can appear so flexible. They aren't memorising one giant solution. They're repeatedly selecting small, reusable pieces of computation as they travel through an internal landscape.

I like to imagine this landscape as an atlas.

An atlas contains many local maps.

No single page describes the whole country in detail. Instead, each page describes one region well. As you travel, you keep turning the pages.

The route you follow determines which map you're looking at.

A neural network can be imagined in much the same way.

Each layer asks, "Which page of the atlas am I on now?"

Having chosen the page, it applies the local knowledge stored there to make the next small step.

Then it asks the question again.

And again.

By the time the network reaches its final layer, it has made dozens or even hundreds of these tiny context-dependent adjustments.

From a distance, the whole process looks like one magical leap from input to output.

Up close, it looks much more like a careful journey, guided by many small steering corrections.

Perhaps that's one reason deep networks are so difficult to understand if we stare at individual layers for too long. Asking what layer 57 "means" is a bit like asking what the 2,413th steering movement meant during a long drive. By itself, not very much.

The important thing is the journey.

Each tiny adjustment makes sense only in the context of all the adjustments that came before and those that will follow.

The intelligence lies not in any single steering correction, but in how thousands of small, context-sensitive decisions work together to reach the destination.